

WHAT IS CLAIMED IS:

1. In a redundant array of inexpensive devices (RAID) comprising a controller and a plurality of storage devices for storing user data, the controller storing a plurality of RAID algorithms to be implemented for writing data to and reading data from the storage devices, a method for RAID conversion comprising:

storing in the controller one or more rules for selecting a desired one of the plurality of RAID algorithms based on one or more conditions of the array;

detecting the one or more conditions of the array;

selecting the desired RAID algorithm based on the detected conditions and the stored rules; and

when the desired RAID algorithm is different from the RAID algorithm currently implemented in the array, converting the array from the currently implemented RAID algorithm to the desired RAID algorithm.

2. The method of claim 1, wherein the converting step comprises:

(a) reading a unit of user data from the storage devices according to the currently implemented RAID algorithm;

(b) defining a watermark indicating the position where the data is read from the current RAID; and

(c) writing user data on appropriate storage devices according to the desired RAID algorithm.

3. The method of claim 2, further comprising:

alternating between performing steps (a), (b) and (c), and processing user I/O requests.

4. The method of claim 1, wherein the desired RAID has fewer storage devices storing user data than the currently implemented RAID.

5. The method of claim 1, wherein the desired RAID has more storage devices storing user data than the currently implemented RAID.
6. The method of claim 1, wherein the conditions include the current capacity utilization of the array.
7. The method of claim 1, wherein the conditions include a performance requirement.
8. The method of claim 1, wherein the conditions include a change in the number of available storage devices in the array.
9. The method of claim 1, wherein the conditions include a decrease in the number of available storage devices in the array.
10. The method of claim 1, wherein the conditions include an increase in the number of available storage devices in the array.
11. The method of claim 1, wherein the conditions include a measure of data criticality of the user data.
12. The method of claim 1, wherein the conditions include a measure of recency and repetition of the user data.
13. The method of claim 1, wherein the conditions include a measure of vulnerability of the storage devices.
14. The method of claim 1, wherein the converting step is performed on line.
15. The method of claim 1, wherein the converting step is performed off line.

16. The method of claim 1, wherein at least some of the RAID algorithms stored in the controller are characterized by a number of storage devices in the array ( $n$ ), and a device-loss insurance level ( $m$ ) such that when up to  $m$  devices of the array are unavailable, user data is fully recoverable from the remaining  $n-m$  devices, where  $1 \leq m < n$ , and wherein the selecting step determines desired  $n$  and  $m$  values based on the detected conditions and the stored rules.
17. The method of claim 16, wherein the device-loss insurance level of the desired RAID is greater than the device-loss insurance level of the currently implemented RAID.
18. The method of claim 16, wherein the device-loss insurance level of the desired RAID is less than the device-loss insurance level of the currently implemented RAID.
19. The method of claim 16, wherein the condition is a decrease in the number of available storage devices in the array and the desired RAID after conversion has the same device-loss insurance level as the currently implemented RAID.
20. The method of claim 16, wherein the condition is a decrease in the number of available storage devices in the array and the desired RAID after conversion has a lower device-loss insurance level than the currently implemented RAID.
21. The method of claim 16, wherein the rules define a maximum device-loss insurance level and a minimum device-loss insurance level for a given  $n$  value, and one or more conditions based on which a desired device-loss insurance level is determined, the desired device-loss insurance level falling between the maximum and minimum device-loss insurance levels.

22. In a redundant array of inexpensive devices (RAID) comprising a controller and a plurality of storage devices for storing user data, the controller storing a plurality of RAID algorithms to be implemented for writing data to and reading data from the storage devices, wherein at least some of the RAID algorithms are characterized by a number of storage devices in the array ( $n$ ), and a device-loss insurance level ( $m$ ) such that when up to  $m$  devices of the array are unavailable, user data is fully recoverable from the remaining  $n-m$  devices, where  $1 \leq m < n$ , a RAID conversion method comprising:

implementing a first RAID algorithm on the array;

selecting a second RAID algorithm characterized by a number of storage devices  $n_2$  and a device-loss insurance level  $m_2$ ,  $n_2$  and  $m_2$  being selectable; and

converting the array from the first RAID algorithm to the second RAID algorithm,

the converting step comprising:

(a) reading a unit of user data from the storage devices according to the first RAID algorithm;

(b) defining a watermark indicating the position where the data is read from the first RAID; and

(c) writing user data on appropriate storage devices according to the second RAID algorithm.

23. The method of claim 22, wherein the writing step includes writing updates to a semi-permanent cache.

24. A redundant array of inexpensive devices (RAID) system comprising:

a plurality of  $n$  storage devices for storing user data thereon; and

a controller connected to the storage devices for controlling writing and reading data to and from the storage devices according to a RAID algorithm,

the controller storing a plurality of RAID algorithms to be implemented for writing data to and reading data from the storage devices,

the controller further storing one or more rules for selecting a desired one of the plurality of RAID algorithms based on one or more conditions of the array,

the controller having stored program instructions or a logic circuit operable to detect the one or more conditions of the array, to select the desired RAID algorithm based on the detected conditions and the stored rules, and when the desired RAID algorithm is different from the RAID algorithm currently implemented in the array, to convert the array from the currently implemented RAID algorithm to the desired RAID algorithm.

25. The system of claim 24, wherein the controller has stored program instructions or a logic circuit operable to convert the array by: (a) reading a unit of user data from the storage devices according to the currently implemented RAID algorithm, (b) defining a watermark indicating the position where the data is read from the current RAID, and (c) writing user data on appropriate storage devices according to the desired RAID algorithm.

26. The system of claim 25, wherein the controller has stored program instructions or a logic circuit operable to alternate between performing steps (a), (b) and (c), and processing user I/O requests.

27. The system of claim 24, wherein the desired RAID has fewer storage devices storing user data than the currently implemented RAID.

28. The system of claim 24, wherein the desired RAID has more storage devices storing user data than the currently implemented RAID.

29. The system of claim 24, wherein the conditions include the current capacity utilization of the array.

- 30. The system of claim 24, wherein the conditions include a performance requirement.
- 31. The system of claim 24, wherein the conditions include a change in the number of available storage devices in the array.
- 32. The system of claim 24, wherein the conditions include a decrease in the number of available storage devices in the array.
- 33. The system of claim 24, wherein the conditions include an increase in the number of available storage devices in the array.
- 34. The system of claim 24, wherein the conditions include a measure of data criticality of the user data.
- 35. The system of claim 24, wherein the conditions include a measure of recency and repetition of the user data.
- 36. The system of claim 24, wherein the conditions include a measure of vulnerability of the storage devices.
- 37. The system of claim 24, wherein the converting step is performed on line.
- 38. The system of claim 24, wherein the converting step is performed off line.
- 39. The system of claim 24, wherein at least some of the RAID algorithms stored in the controller are characterized by a number of storage devices in the array (n), and a device-loss insurance level (m) such that when up to m devices of the array are unavailable, user data is fully recoverable from the remaining n-m devices,

where  $1 \leq m < n$ , and wherein the selecting step determines desired  $n$  and  $m$  values based on the detected conditions and the stored rules.

40. The system of claim 39, wherein the device-loss insurance level of the desired RAID is greater than the device-loss insurance level of the currently implemented RAID.

41. The system of claim 39, wherein the device-loss insurance level of the desired RAID is less than the device-loss insurance level of the currently implemented RAID.

42. The system of claim 39, wherein the condition is a decrease in the number of available storage devices in the array and the desired RAID after conversion has the same device-loss insurance level as the currently implemented RAID.

43. The system of claim 39, wherein the condition is a decrease in the number of available storage devices in the array and the desired RAID after conversion has a lower device-loss insurance level than the currently implemented RAID.

44. The system of claim 39, wherein the rules define a maximum device-loss insurance level and a minimum device-loss insurance level for a given  $n$  value, and one or more conditions based on which a desired device-loss insurance level is determined, the desired device-loss insurance level falling between the maximum and minimum device-loss insurance levels.

45. A computer program product comprising a computer usable medium having a computer readable code embodied therein for controlling a redundant array of inexpensive devices (RAID), the RAID comprising a controller and a plurality of storage devices for storing user data, the controller storing a plurality of RAID

algorithms to be implemented for writing data to and reading data from the storage devices, the computer program product comprising:

first computer readable program code configured to cause the controller to storing one or more rules for selecting a desired one of the plurality of RAID algorithms based on one or more conditions of the array;

second computer readable program code configured to cause the controller to detect the one or more conditions of the array;

third computer readable program code configured to cause the controller to select the desired RAID algorithm based on the detected conditions and the stored rules; and

fourth computer readable program code configured to cause the controller to, when the desired RAID algorithm is different from the RAID algorithm currently implemented in the array, convert the array from the currently implemented RAID algorithm to the desired RAID algorithm.

46. The computer program product of claim 45, wherein the fourth computer readable program code comprises:

fifth computer readable program code configured to cause the controller to read a unit of user data from the storage devices according to the currently implemented RAID algorithm;

sixth computer readable program code configured to cause the controller to define a watermark indicating the position where the data is read from the current RAID; and

seventh computer readable program code configured to cause the controller to write user data on appropriate storage devices according to the desired RAID algorithm.

47. The computer program product of claim 46, further comprising

seventh computer readable program code configured to cause the controller to process user I/O requests; and



eighth computer readable program code configured to cause the controller to alternate between executing the fifth, sixth and seventh program codes and executing the seventh program code.

48. The computer program product of claim 45, wherein the desired RAID has fewer storage devices storing user data than the currently implemented RAID.

49. The computer program product of claim 45, wherein the desired RAID has more storage devices storing user data than the currently implemented RAID.

50. The computer program product of claim 45, wherein the conditions include the current capacity utilization of the array.

51. The computer program product of claim 45, wherein the conditions include a performance requirement.

52. The computer program product of claim 45, wherein the conditions include a change in the number of available storage devices in the array.

53. The computer program product of claim 45, wherein the conditions include a decrease in the number of available storage devices in the array.

54. The computer program product of claim 45, wherein the conditions include an increase in the number of available storage devices in the array.

55. The computer program product of claim 45, wherein the conditions include a measure of data criticality of the user data.

56. The computer program product of claim 45, wherein the conditions include a measure of recency and repetition of the user data.

57. The computer program product of claim 45, wherein the conditions include a measure of vulnerability of the storage devices.

58. The computer program product of claim 45, wherein the converting step is performed on line.

59. The computer program product of claim 45, wherein the converting step is performed off line.

60. The computer program product of claim 45, wherein at least some of the RAID algorithms stored in the controller are characterized by a number of storage devices in the array ( $n$ ), and a device-loss insurance level ( $m$ ) such that when up to  $m$  devices of the array are unavailable, user data is fully recoverable from the remaining  $n-m$  devices, where  $1 \leq m < n$ , and wherein the selecting step determines desired  $n$  and  $m$  values based on the detected conditions and the stored rules.

61. The computer program product of claim 60, wherein the device-loss insurance level of the desired RAID is greater than the device-loss insurance level of the currently implemented RAID.

62. The computer program product of claim 60, wherein the device-loss insurance level of the desired RAID is less than the device-loss insurance level of the currently implemented RAID.

63. The computer program product of claim 60, wherein the condition is a decrease in the number of available storage devices in the array and the desired RAID after conversion has the same device-loss insurance level as the currently implemented RAID.

64. The computer program product of claim 60, wherein the condition is a decrease in the number of available storage devices in the array and the desired RAID after conversion has a lower device-loss insurance level than the currently implemented RAID.

65. The computer program product of claim 60, wherein the rules define a maximum device-loss insurance level and a minimum device-loss insurance level for a given n value, and one or more conditions based on which a desired device-loss insurance level is determined, the desired device-loss insurance level falling between the maximum and minimum device-loss insurance levels.